Baking - Creating blocks on Tezos

Contents

Introduction

Introduction

The following characteristics set the Tezos blockchain apart from other blockchains:

1. **on-chain governance**, making it possible for the protocol to evolve by upgrading itself through successive amendments
2. the **liquid proof of stake** (LPOS) consensus
3. the **Michelson development language** used for smart contracts, providing for verification based on formal evidence

This document provides a detailed introduction

- to the consensus algorithm
- to the mechanism for creating blocks — or *baking*, as it is known on Tezos.

Definition of the main concepts

Definition of the main concepts 1/3

- **Block:** A sequence of operations for validation. A block contains information such as the current block number, information about the previous block, and also the time at which it was created.
- **Baking:** Creation of new blocks on the Tezos blockchain by **bakers**, who are given a reward for each block that they bake. A baker may also be referred to as a **delegate**.
- **Endorsement:** In addition to creating new blocks, bakers may approve blocks made by other bakers. They are known as **endorsers** of the block and are rewarded for doing so.
- **Delegation:** Every tez (XTZ or ꜩ) holder can transfer the *baking and endorsing rights* and *voting rights* associated with their tez to a baker. The holder always retains control of their funds and may change delegate at any time.

Definition of the main concepts 2/3

- **Roll:** One roll is equal to 8000ꜩ. To have the right to be recognized as a baker in a network, a tez holder must possess at least one roll in their staking balance. oting rights are also weighted by the number of rolls (1 roll = 1 vote = 8000ꜩ).
- **Cycle:** Unit of time during which 4,096 blocks are created on the Tezos blockchain. It corresponds to roughly 2 days, 20 hours, and 16 minutes (based on 1 minute per block created if no baker has failed).
- **Fitness:** Score representing the quality of the chain up to a given block. In the current protocol, the fitness level equals the block height in the chain.
- **Double baking** and **double endorsing** are defined on slide 35.
- **Slot:** A slot is a position available for a block or an endorsement: each cycle has 4,096 baking slots and 131,072 endorsement slots. At the start of each cycle the protocol assigns a list of bakers to each slot and exactly one baker to each endorsing slot using a uniform law weighted by the number of rolls. The first baker in the baking list is tasked with baking. If they don't, the next baker is given the task, etc.

Definition of the main concepts 3/3

- **Node**[1]**:** A Tezos node is a peer (a machine) on the peer-to-peer network. It keeps a copy of the current snapshot and propagates the blocks baked and operations performed to the other peers.
  A node is not necessarily a baker, but a baker is always associated with one or more nodes. Nodes are usually implemented in one of the following three ways:
  - In **archive** mode and containing the whole history of the chain. As of year end 2020, an archive node takes up around 80GB.
  - In **rolling** mode and containing the minimum data required for to operate the chain, with the oldest information being deleted regularly.
  - In **full** mode (the default mode). Full mode combines elements of archive and rolling modes: it contains all the information in a rolling node, plus the minimum data required to reconstitute the entire chain from the genesis block.

  Other types of node could also be created.

[1] Click here for more information about nodes

Contents of a block

Contents of a block

A *blockchain* is made up of a series of *blocks*. Each block consists of a *header*, and a list of *operations*. The header contains key data related to the block's validity, such as:

- the block level, or the block height in the chain
- the block predecessor's hash, which is used to chain the blocks together
- the block's timestamp
- the signature of the baker that produced the block
- the block's fitness, assigning a score to the chain up to the current block

The block's body then contains a list of various operations (transfer of tez, deployment of smart contracts, etc.) up to the maximum size permitted under the protocol (currently 512KB per block).

Liquid proof of stake and delegation

Liquid proof of stake and delegation

Baking rights are attributed on Tezos randomly based on the number of rolls held (1 roll = 8000ꜩ).

A tez holder may:

- *become a baker* if they possess at least one roll in own funds or by delegation
- *delegate their tez to a baker*, to participate in the consensus

In practice, certain bakers may pass on some of the rewards they receive to the individuals who delegated their tez to them. This distribution may encourage others to delegate their tez, thereby increasing the number of rolls held by these bakers and thus their chance of baking/endorsing.

Liquid proof of stake (LPOS) provides the option for a tez holder to retain ownership of their funds while delegating their associated (voting, baking) rights to a baker. This principle differs from the delegated proof of stake (DPOS) of other blockchains, in which participants vote for a restricted number of block validators.

Baking

Baking

In Tezos we call participants in the block validation process **delegates**. Delegates can

- act as a *baker*, creating and signing blocks, or
- act as an *endorser*, approving blocks by issuing an endorsement operation (endorsements visible at block level $n + 1$ relate to block level n).

Baking and endorsement rights are attributed at random, several cycles in advance. To deter dishonest behavior (e.g. double baking or double endorsement), **each delegate must place a fixed quantity of tez as a security deposit for a limited time** (five cycles), which will be confiscated if they attempt to compromise the chain (double baking/endorsement).

Prospective bakers must meet the following requirements:

- server available round the clock and stable internet connection
- at least 8GB of RAM
- SSD disk (preferably with more than 100GB storage)
- at least one roll (8000ꜩ)

Endorsement

Endorsement

A list of endorsers is drawn up for each block.

Like bakers, endorsers are selected at random from among the existing delegates.

To approve a block at level n, the endorsers each add a signature operation that is included in the $n+1$-level block. Once block $n+1$ is baked, no more endorsement operations are accepted for block n.

Each block has 32 endorsement slots. There is no obligation for a baker to include 32 endorsements but their reward is proportional to the number of endorsements included.

An endorser may obtain several endorsement slots for a given block. They will then be rewarded for each endorsement completed (i.e. for each slot validated).

Lastly, just as for bakers, endorsers must provide a security deposit to curb the risks of dishonest behavior compromising the chain (i.e. double endorsement).

Options for baking

Options for baking

In practice, blocks are baked and endorsed automatically. There are several ways of baking, each with its own advantages and disadvantages. Here are some baking options, with links to more information:

- Kiln and Ledger Nano S:
  `https://gitlab.com/tezos-kiln/kiln`
  An intuitive plug-and-use method for setting up a node and baker.
- Remote Signer and Ledger Nano S:
  `https://github.com/obsidiansystems/ledger-app-tezos`
- Remote Signer via the Cloud:
  `https://www.ecadlabs.com/signatory`

The following links provide information on how to set up a baker:

- `https://github.com/tzConnectBerlin/baking-support`
- Tezos Node Setup & Maintenance by BakingBenjamins

Remote signing

Remote signing 1/2

Remote signing software is one way to reduce the exposure of private keys during baking operations. In the block validation process (baking/endorsement), the software interacts with a private key management system (e.g. Ledger Nano S, a remote machine communicating via a secure channel).

Remote signing involves physically separating baking/endorsement requests originating from the network and intended for a given delegate, from that delegate's signing operation. The Tezos client has already implemented signing mechanisms (unix, tcp, http, and https) that can be used to send signing requests via the channel of communication chosen by the baker.

Remote signing 2/2

The following events occur during a remote signing operation:

1. The network sends a signing request to the remote signing software that interacts with the private key management system.
2. The remote signing software verifies what the operation involves — usually *creation of a block* or an *endorsement*[1] — then sends it to the key management system, which signs the operation.
3. The signature is then validated by the remote signing software and sent back to the Tezos node, which propagates the operation to the network.

The key should preferably be **held in a hardware vault** so it is never directly accessible on the network. That said, it is more secure to ensure that the hardware vault authorizes only block signing rather than the signing of just any transaction (i.e. so as not to create a blind signing mechanism).

For example: Ledger Nano S together with Kiln meets these two conditions.

[1] The third, rarer option is a *nonce revelation*.

Baker selection

Baker selection

In each cycle, a random seed is generated based on certain information contained in the previous blocks.

The seed for cycle n is used to select a **roll snapshot** (a picture of the status of the rolls of the participants in the network) in cycle $n - 2$.
The seed can then be used to select the rolls (from those in the snapshot) to which baking/endorsement rights are to be attributed for cycle $n + 5$.

Rights are allocated using the following mechanism:

- For each block, an infinite list is drawn up, setting an order of priority for bakers and endorsers.

Once listed, a baker therefore has to wait for at least seven cycles before being eligible for selection.

Rewards for bakers and endorsers

Rewards for bakers and endorsers

Bakers are crucial because they operate the chain by creating blocks and checking the integrity of the data they contain. Each block baked leads to the creation of new tez to incentivize the baker and the endorsers to perform this task.

The baker's reward depends on their level of priority and the number of endorsements gained for the block they are baking, using the formula:

$$\text{baking reward} = \begin{cases} e * 1.25 & \text{if priority} = 0 \\ e * 0.1875 & \text{otherwise} \end{cases} \quad \text{where } e = \text{number of endorsements}$$

i.e. a maximum of 40ꜩ for the baker, not counting the transaction fees included in their block.
The reward for endorsement operations also depends on their level of priority based on the formula:

$$\text{endorsing reward} = \begin{cases} e * 1.25 & \text{if priority} = 0 \\ e * 0.833333 & \text{otherwise} \end{cases} \quad \text{where } e = \text{number of endorser's slots}$$

An endorser may have several of the 32 slots available on a given block.

Issuing new tez

Issuing new tez 1/2

Each time a block is created and validated, new tez are created.

The per-block reward can be used to estimate that a block produces a maximum of 80ꜩ.
Extending this per-block reward to the number of blocks baked per annum at a rate of one block per minute (52,560 per annum excepting leap years), and assuming all rolls are active (19,077 in block 0), then the number of tez increases by at most 5.51% in year 1 (in 2018).
Since the rewards are fixed, this percentage will decline over the years.

The growth in the ``money supply'' does not apply solely to bakers. If a tez holder delegates them to a baker who passes on a portion of their rewards, that tez holder will also benefit from the new money created.

Issuing new tez 2/2

Several phenomena independent of the baker can lead to fluctuations in their return:

- Not all rolls might be active in the baking process, increasing annual return.
- The time interval between blocks could be more than a minute, decreasing annual return.
- The full 32 endorsements might not have been gained, decreasing annual return.

A holder delegating their tez may wish to select a reliable baker who pays them interest. The estimated return can be observed based on the amount of tez delegated:

- `https://www.stakingrewards.com/earn/tezos`

Modeling a baker's gains

Modeling a baker's gains 1/2

With the following assumptions …

1. Number of active rolls in year one = 84,156 (correct at september 2020)
2. One block baked per minute
3. 32 endorsements for each block

…a baker possessing one roll (= 8000ꜩ) will create on average one block around every two months, and will endorse around 16 blocks per month. This is close to 480ꜩ earned per annum on average, or a return of approximately 6%. This value is slightly higher than the 5.51% stated above because we assume here that there are inactive rolls, which means active rolls earn more revenue than in a situation where all rolls are active.

A baker also receives fees for the transactions they include in any blocks they create. These fees depend on supply and demand, but the gains are currently well below the reward for creating the block itself (≈ 0.1% of the baker's gains in 2019).

Modeling a baker's gains 2/2

Continuing the assumptions from slide 30, here is an example showing what a baker who retains their baking rewards can earn in a year, according to how many tez they hold:

Year	Total kꜩ in circulation	Alice			Bob		
		ꜩ	rolls	ROI	ꜩ	rolls	ROI
1	673248	8,000.00ꜩ	1.00	6.25%	80,000.00ꜩ	10.00	6.25%
2	715296	8,499.64ꜩ	1.06	5.53%	84,996.44ꜩ	10.62	5.53%
3	757344	8,969.92ꜩ	1.12	4.95%	89,699.16ꜩ	11.21	5.45%
4	799392	9,414.08ꜩ	1.18	4.47%	94,584.95ꜩ	11.82	4.89%
5	841440	9,834.88ꜩ	1.23	4.06%	99,213.75ꜩ	12.40	4.84%
6	883488	10,234.65ꜩ	1.28	3.72%	104,011.01ꜩ	13.00	4.76%
7	925536	10,615.40ꜩ	1.33	3.42%	108,960.70ꜩ	13.62	4.34%
8	967584	10,978.84ꜩ	1.37	3.17%	113,685.52ꜩ	14.21	4.28%
9	1009632	11,326.50ꜩ	1.42	2.94%	118,552.67ꜩ	14.82	3.93%
10	1051680	11,659.67ꜩ	1.46	2.74%	123,217.12ꜩ	15.40	3.89%

Alice has only 1 roll at the outset and Bob has 10. After ten years, Alice is still baking with a single roll whereas Bob has five more. This reveals the **critical mass** applicable to bakers. Unlike smaller bakers, larger ones earn a linear increase in the number of rolls they hold over time.

These results are simulations based on assumptions, and so a baker's actual revenue may not be the same as the figures shown in this table.

Security deposits and penalties

Security deposits and penalties

To discourage double baking/endorsement[1] **a security deposit is required**:

- 512ꜩ per block baked
- 64ꜩ per endorsement slot

The security deposit remains physically in the account but is locked up for a period of five cycles (around two weeks). The corresponding tez are always taken into account when calculating slots, but they cannot be spent or used as a new security deposit until they have been released again.

Bakers are advised to hold at least 10% of the total number of tez used for baking purposes. If so, they will comfortably be able to provide the security deposits required for each baking/endorsement operation. Otherwise, they may find it impossible to place the security deposit and thus miss baking/endorsement operations. That incurs an opportunity cost, but does not incur any penalty charge.
Where double baking/endorsement is found to have taken place, an accusing baker may provide evidence in a block for a period of five cycles starting from the fake operation. As a penalty, half the culpable baker's deposit goes to the accuser, and the other half is destroyed.

[1] *Double baking* (or *double endorsement*) means baking (endorsing) two different blocks at the same level with the same key, which may temporarily cause the chain to fork.

Actions prejudicial to the network

Actions prejudicial to the network

The following actions result in the punitive seizure of the security deposit:

- **Double baking:** A dishonest baker may try to bake two blocks at the same block height in the chain. However, for them to be able to do this, they would have be drawn at random and then deposit 512ꜩ per block baked.

- **Double endorsement:** A dishonest endorser could attempt to endorse two blocks at the same block height in the chain. For this to happen, the endorser would have to be drawn at random and then place a security deposit of 64ꜩ per endorsement slot.

- **Seed slash:** a baker is asked to provide random seeds to feed into the baker selection algorithm. To avoid the algorithm being deterministic, the seeds are given in hashed form. The baker undertakes to deliver in the following cycle the seeds corresponding to the hashes sent. If they do not do so, they do not earn a reward but may still get back the associated deposit.

The protocol allows a window of five cycles during which an **accuser** may demonstrate evidence of double baking or double endorsement (several times if necessary). The dishonest baker will then lose the security deposit(s) placed for the fake operation up to the latest point in the cycle (which may be costly, especially if the baker has done a lot of baking/endorsing over the relevant period).

Minimum time delay before baking a block

Minimum time delay before baking a block

The minimum time delay between two blocks is 1 minute. The priority 0 baker may then bake at least 60 seconds after the final block. The priority 1 baker will have to wait for at least 40 more seconds, the priority 2 baker for another 40 seconds, and so on.

To be able to bake as rapidly as possible, the baker must gain enough endorsements (at least 24 of 32). Otherwise, they must observe a time penalty of several seconds per missing endorsement based on the following formula:

$$\text{Time delay} = 60 + 40 * p + 8 * \max(0,\ 24 - e) \text{ seconds} \quad \text{where} \quad \begin{array}{l} p = \text{baker's priority} \\ e = \text{number of endorsements} \end{array}$$

Any block baked ahead of the minimum time delay is automatically rejected under the protocol's rules.

Over-delegation

Over-delegation

A security deposit is required for each baking/endorsement operation. This amount remains locked up for five cycles. In addition, a baker cannot spend tez that have been delegated to them or use them as security deposits. Consequently, a baker must possess sufficient funds to bake/endorse blocks.

A baker is said to be over-delegated when they do not have sufficient own funds relative to their number of rolls.

The baker will be unable to provide the security deposit required for all the blocks they have been awarded. The baker will therefore be unavailable for certain blocks, causing the chain to slow down because it will have to wait for a lower-priority baker.

Measures are being considered that would allow bakers to refuse delegation and avoid finding themselves in an over-delegated position. This could become possible if a future Tezos amendment including these measures is passed.

Inactivity

Inactivity

A baker who for five cycles refuses to bake or endorse or for technical reasons is unable to do so, will be treated as inactive by the chain, and can no longer be picked as a baker/endorser.

The inactive baker must wait seven cycles after reactivating their account, before they can bake again.

This mechanism is designed to ensure that a baker who is no longer active, does not slow down the chain.

Attacks

Attacks

The attacks potentially targeting a baker include:

- **DoS (denial of service attack):** Suppose Bob knows that Alice is about to bake a block with priority 0. He may decide to launch a DoS attack to prevent Alice from publishing her block correctly, thus damaging Alice's reputation. If the attack is sustained, and if Bob appears toward the top of the list of potential bakers (especially if he has priority 1), then he may also be able to steal Alice's block.

- **Theft or loss of a private key:** Holding the private key is the only way to gain access to the associated tez and to sign various operations. Thus it is vital to keep the private key safe and secure (e.g. in a hardware wallet).

- **Eclipse attack:** An attacker can try to cut off (or take control of) the connection between the baker and the network nodes to which they are connected. The result is similar to a DoS attack, but the attacker can also send the isolated baker blocks that are not part of the consensus.

Summary diagram

nomadic labs
Summary diagram

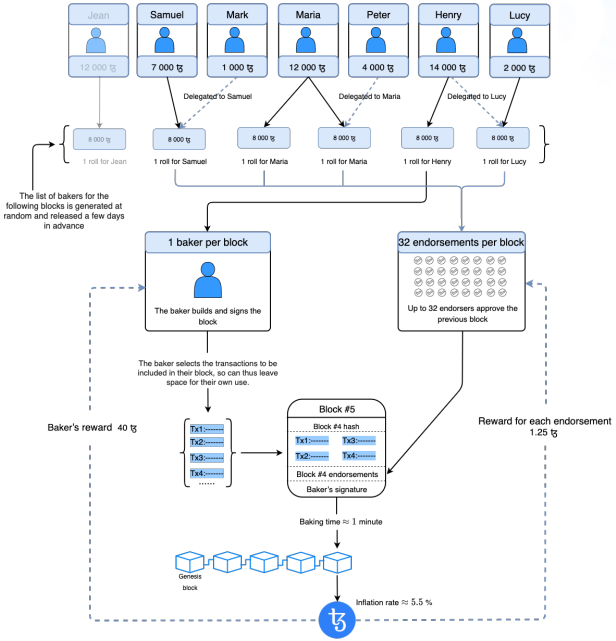

nomadic labs

Let's keep the conversation going

https://tezos.com
https://developers.tezos.com
https://tezos-baking.slack.com